

REMARKS

This Preliminary Amendment is filed in connection with a Request for Continued Examination and a 2-month extension of time, and is in response to the Final Office Action mailed Feb. 7th, 2005 and the Advisory Action mailed April 29th, 2004. All objections and rejections are respectfully traversed.

Claims 1-85 and 87-104 are now in the case.

Claim 86 has been cancelled without prejudice.

New claims 88-98 have been added.

Claims 20, 43, 65, and 79 have been amended to correct minor typographical errors.

Election/Restrictions

At paragraph 2 of the Final Office Action, the Examiner restricted claim 86 from the other claims, stating it was directed to an independent and distinct invention. The Applicant has cancelled claim 86 to advance the prosecution of this case.

Claim Rejections – 35 U.S.C. §103

At paragraph 2 of the Final Office Action, claims 1-85 and 87 were rejected under 35 U.S.C. 103(a) as being unpatentable over Han et al., U.S. Patent No. 6,158,017, issued on Dec. 5th, 2000 (hereinafter Han).

The Applicant's invention, as set forth in representative claim 1, comprises in part:

1. A method for enabling recovery from two or fewer concurrent failures of storage devices in a storage array, the method comprising the steps of:
providing the array with a predetermined number of storage devices, including *a plurality of first devices configured to store data and*

row parity, and one diagonal parity device configured to store diagonal parity, wherein the predetermined number of storage devices n is $p+1$ and wherein p is a prime number;

dividing each device into blocks;

organizing the blocks into stripes that contain a same number of blocks in each device, wherein each stripe comprises $n-2$ rows of blocks;

defining the diagonal parity along diagonal parity sets that span the first devices, wherein the diagonal parity sets wrap around within a group of $n-2$ rows so that all blocks belonging to diagonal parity sets of a stripe are stored in the stripe; and

computing and storing the diagonal parity for all of the diagonal parity sets except one on the diagonal parity device.

Han discloses two different parity arrangements, DH and DH2.

Han's DH arrangement involves a disk array of N disks (N = a prime number) with each disk logically divided into $N-1$ blocks to form a matrix of dimension $(N-1)*N$. *See col. 5, lines 21-25 and Fig 3.* The data blocks of each row of the array are defined as a horizontal parity group, and horizontal parity is stored on a diagonal line of blocks striped across disks of the array. *See col. 5, lines 50-60 and Fig 3.* Diagonal parity is computed along diagonal lines of blocks, and stored on a row of blocks striped across the disks of the array. *See col. 6, lines 12-30 and Fig 3.* One diagonal parity block is computed entirely from horizontal parity blocks. *See col. 7, lines 33-66.*

Han's DH2 arrangement involves a disk array of $N+1$ disks (N = a prime number), with each disk logically divided into $N-1$ rows of blocks to form a matrix of dimension $(N-1)*(N+1)$. *See col. 10, lines 64-66 and Fig. 4.* The data blocks in each row of disks (0 to N) are defined as a horizontal parity group, and horizontal parity is stored on the last disk (disk $N+1$) of the array. *See col. 11, lines 13-22.* Further, diagonal parity is computed along diagonal lines of blocks in the array. These diagonal lines traverse the data disk (disks 0 to N) of the array, but excluding the horizontal parity disk (disk $N+1$).

See col. 11, lines 36-45 and Fig. 4. The diagonal parity is stored on a row of blocks striped across the data disks of the array *See* col. 12, line 65 to col. 13 line 14 and Fig 4.

The Applicant respectfully urges that Han is silent concerning the Applicant's claimed invention relating to "*a plurality of first devices configured to store data and row parity,*" and "*one diagonal parity device configured to store diagonal parity,*" and "*defining the diagonal parity along diagonal parity sets that span the first devices,*" and "*computing and storing the diagonal parity for all of the diagonal parity sets except one on the diagonal parity device.*"

While the Applicant teaches "*one diagonal parity device configured to store diagonal parity*", Han discloses, in both his DH and DH2 arrangements, striping diagonal parity across multiple disks of an array. *See* Fig 3. (diagonal parity blocks across row b3) and Fig 4. (diagonal parity blocks across row b3). As discussed in detail below, storing diagonal parity on a single disk has advantages over striping. Further, the Applicant respectfully disagrees with the Examiner's assertion that one of ordinary skill could simply "rearrange" Han to show the Applicant's novel invention.

Storing diagonal parity on *one* device provides advantages over other configurations. For example, the Applicant comments at page 26, lines 13-18 of the Application:

Advantageously, the invention provides double disk failure parity protection in, e.g., a RAID-4 style concentrated parity format, where all the parity information is stored on two devices, such as disks. The inventive parity technique thus enables data disks to be incrementally added to disk arrays without reformatting or recalculating the existing parity information. The invention uses a minimum amount of redundant disk space, i.e., exactly two disks per array.

In contrast, Han's striped diagonal parity lacks these features, for example, requiring considerable reconfiguration if disks are added and using more than the minimum amount of redundant disk space (topics discussed in detail below).

Further, one of ordinary skill could not simply "rearrange" Han to show the Applicant's novel invention. In the Office Action of Feb, 7th 2005, the Examiner asserted at page 5 (emphasis added):

Regarding independent Claims 1, 7, 18, 24, 30, 36, 41, 42, 64, 87, *Han does not disclose the feature of 'one diagonal parity device configured to store diagonal parity'*. In Fig 10 Han depicts ... *It would have been obvious to a person having ordinary skill in the art at the time the invention was made to rearrange the matrix of Han by incorporating an additional dedicated disk* for storing the diagonal parity group since the matrix is expandable according to N+1 disks, where if N=7, then the number of disks is N+1=8.

The Applicant respectfully disagrees with these statements, and asserts such a "rearrangement" is taught away from by Han and would be impossible barring significant departure from the teaching of Han and considerable innovation.¹

First, Han teaches away from using only one disk for diagonal parity and disapproves such an approach. In reference to the EVENODD technique that stores diagonal parity on a single disk Han comments disapprovingly at col. 1, line 62 to col. 2, line 9 (emphasis added):

This technique suffers from two drawbacks, one of which *bottlenecks the parity disks because of their undistributed state* ... Namely, if m (a prime number) disk are used as data disks, only two parity disks may be added, thus minimizing parity disk overhead. This technique, however, *may de-*

¹ The Applicant further disagrees with the Examiner's assertion that "the matrix is expandable according to N+1 disks. The embodiment of Fig. 10 to which the Examiner refers already has N+1 disks. Thus, if one were to add an additional disk for diagonal parity there would be N+2 disks.

grade performance due to bottlenecks as described in the previous two dimensional technique...

Han goes on to consistently teach diagonal parity stripped across multiple disks. While, Han varies the location of horizontal parity, all of Han's embodiments teach stripping diagonal parity. Thus, one of ordinary skill reading Han would be dissuaded from placing diagonal parity on a single storage device.

Second, even if one of ordinary skill were inclined to attempt to "rearrange" Han by placing diagonal parity on a dedicated disk, this could not be accomplished without significant departure from the teachings of Han and further innovation. The Examiner suggests one could simply add an extra disk to the DH2 parity arrangement shown in Fig 10. Yet, Fig. 10 shows a stripe of 7 diagonal parity blocks (labeled D0, D1, D2, D3, D4, D5, D6) while a single disk in Han has only has 6 blocks (See rows labeled b0, b1, b2, b3, b4, b5). The 7 diagonal parity blocks could not possibly fit on a 6 block disk.

Stated more generally, Han's teaches a DH2 parity arrangement with $N+1$ disks, each disk with $N-1$ blocks, and where N diagonal parity blocks are calculated. The N diagonal parity block would not fit on an $N-1$ block disk. One would always have an extra block and that would have to be stored somewhere else.² Given this problem, it would not have been obvious to rearrange Han in the manner the Examiner suggests as a significant obstacle would have to be overcome.

The Applicant discusses this obstacle in the background section of the Application in reference to the EVENODD technique. The EVENODD technique attempted to

² Similarly, one could not "rearrange" Han's DH parity scheme for the same reason. Looking to Fig 6, there are 6 diagonal parity blocks (D0,D1,D2,D3,D4,D5) and only 5 blocks of storage space per disk. Thus simply adding another disk would not permit diagonal parity to be stored on a single device.

solve the problem by XORing the extra diagonal parity block into the other diagonal parity blocks. The Applicant describes at page 6, lines 11-15:

Therefore, the parity of the extra (missing) diagonal parity block (X) is recorded by XOR'ing that diagonal parity into the parity of each of the other diagonal parity blocks. Specifically, the parity of the missing diagonal parity set is XOR'd into each of the diagonal parity blocks P4 through P7 such that those blocks are denoted P4X-P7X.

This solution unfortunately required additional computations and generally decreased overall system performance making it unsuitable for many applications.

Accordingly, the Applicant's claimed invention improves upon both Han and other prior techniques by teaching "*a plurality of first devices configured to store data and row parity*," "*defining the diagonal parity along diagonal parity sets that span the first devices* " and "*computing and storing the diagonal parity for all of the diagonal parity sets except one on the diagonal parity device.*" Using this novel arrangement, the Applicant can recover from any two concurrent disk failures despite missing one diagonal parity block.

Accordingly, the Applicant respectfully urges that Han is legally insufficient to make obvious the presently claimed invention under 35 U.S.C. § 103 because of the absence of the Applicant's claimed novel "*a plurality of first devices configured to store data and row parity*," and "*one diagonal parity device configured to store diagonal parity*," and "*defining the diagonal parity along diagonal parity sets that span the first devices*," and "*computing and storing the diagonal parity for all of the diagonal parity sets except one on the diagonal parity device.*"

New Claims

The Applicant has added new claims 88-90 to the Application. To advance the prosecution of the case, the Applicant would like to preemptively discuss these claims in relation to Han. The Applicant respectfully directs the Examiner's attention to page 25, lines 9-16 for support for these new claims.

Applicant's invention, as set forth in new representative claim 88, comprises in part:

88. (NEW) A method for enabling recovery from two or fewer concurrent failures of storage devices in a storage array, the method comprising the steps of:

associating the array with a predetermined maximum number of storage devices, wherein the predetermined maximum number of storage devices n is $p+1$ and wherein p is a prime number;

providing less than p first storage devices configured to store data and row parity;

providing a diagonal parity storage device configured to store diagonal parity;

configuring the array to treat one or more absent storage devices as a present storage device that contains zero-valued data;

dividing the storage devices into blocks;

defining the diagonal parity along diagonal parity sets that span the first storage devices and the one or more absent storage devices, wherein the diagonal parity sets wrap around within a group of $n-2$ rows; and

computing the diagonal parity for all of the diagonal parity sets except one, and storing the diagonal parity on the diagonal parity device.

The Applicant novelly claims "*providing less than p first storage devices configured to store data and row parity*" and "*configuring the arrays to treat one or more absent storage devices as a present storage devices that contains zero-valued data.*" In

sharp contrast, Han requires each and every device of his N or $N+1$ storage devices (depending on the embodiment) to be physically present. Since Han stores diagonal parity data on a stripe across the disks of the array, every disk is required to store its share of parity information. This is a significant limitation since many real-world users desire to start out with a small number of storage devices, and add storage space as needed. Accordingly, the Applicant provides flexibility by teaching an array that may operate with less than p first storage devices present.

In the event that the Examiner deems personal contact desirable in disposition of this case, the Examiner is encouraged to call the undersigned attorney at (617) 951-3078.

All independent claims are believed to be in condition for allowance.

All dependent claims are believed to be dependent from allowable independent claims.

Applicant respectfully solicits favorable action.

Please charge any additional fee occasioned by this paper to our Deposit Account No. 03-1237.

Respectfully submitted,



James A. Blanchette
Reg. No. 51,477
CESARI AND MCKENNA, LLP
88 Black Falcon Avenue
Boston, MA 02210-2414
(617) 951-2500